Weitek Announces SPARC Upgrade Chip

80-MHz CPU Gives 40%-50% Performance Boost to SPARCstation 2

By Linley Gwennap

Taking a page from Intel's OverDrive strategy, Weitek has announced a clock-doubled upgrade processor for SPARCstation 2 systems. The 80-MHz "SPARC Power" chip provides about a 50% increase in performance over systems that use Weitek's W8701 CPU (or Fujitsu's pin-compatible MB86903). Weitek estimates that the installed base of such systems—including Sun's SPARCstation 2 (SS2) and IPX workstations and third-party clones—is over 250,000 units. These users will be able to purchase upgrade kits, shipping by the end of this quarter, for $1500 in single-piece quantities.

Although Sun did not design its products with an upgrade socket such as those in some PCs, the company has traditionally socketed all of its CPUs to simplify replacing defective parts. Sun's socket strategy left the door open for the possibility of an upgrade processor, but Sun itself chose not to develop such a chip, preferring to push its customers toward the next-generation SPARCstation 10 (SS10) products.

Weitek, which has lost its lucrative role as one of Sun's primary CPU suppliers to Texas Instruments (TI), has leapt into this breach with SPARC Power (SPX2 for short), which could prop up its SPARC revenues for another year or two. Although the volume of upgrade sales will be lower than Sun's chip purchases, Weitek will market the SPX2 directly to end users, resulting in much higher profit margins.

Perversely, Sun's current software strategy also encourages SS2 and IPX users to upgrade their processors rather than switch systems. Many users are upset that Sun's Solaris 2.0 operating system breaks binary compatibility with SunOS applications and also feel that Solaris is significantly slower than the older SunOS. Thus, users who bought SunOS with their SS2 or IPX are faced with purchasing a completely new hardware platform and a new OS, plus obtaining new binaries for all of their software, just to move to the next performance level. The Weitek upgrade allows these users to stay with their current hardware and software, gaining a significant performance advantage simply by swapping CPU chips.

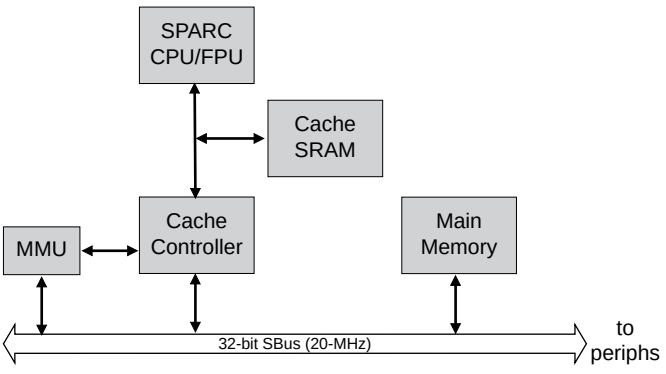

Figure 1. A simplified diagram of a SPARCstation 2 or IPX system, which can accommodate either the original W8701 CPU or the new SPARC Power upgrade processor.

The $1500 upgrade kit includes a SPARC Power processor, a specially-designed chip puller/installation tool, an anti-static wrist strap, and even a small flashlight to aid in the installation process, along with detailed instructions. The upgrade process takes only a few minutes and can be performed by most technically-inclined end users.

Weitek plans to sell the upgrade kit through its direct sales force in North America but will use distributors in Europe and Japan. Although Sun has "endorsed" the upgrade product, Weitek has selected European and Japanese distributors that do not market Sun products, thus avoiding channel conflicts.

Socket-Compatible Upgrade

Weitek describes its new upgrade chip as "socket-compatible" with its older 8701 (see µPR 5/1/91, p. 13). It is fully software-compatible with SPARC version 7 and can be used in any system design that uses the LSI 6484x system-logic chip set. The SPX2 is not strictly pin-compatible with the 8701 because some functions not used with the LSI chip set have been left out of the SPX2 to simplify the design.

The system design of the SS2, shown in Figure 1, creates several problems for an upgrade processor. The 40-MHz 8701 is badly hampered by the single 32-bit path to cache memory. Simply increasing the CPU clock to 80 MHz would have exacerbated this memory bottleneck. The designers chose to add data and instruction caches on the processor chip, reducing the number of memory references that go off-chip. These on-chip caches had to be carefully designed, however, to be totally transparent to the external cache, which was not originally intended to be a second-level cache.

Like other field-upgrade processors, the SPX2 uses a phase-locked loop (PLL) circuit to double the frequency of the system clock. Thus, the new design must work at exactly 80 MHz; no slower or faster will do. The upgrade chip also has to operate within the thermal constraints of the original system.

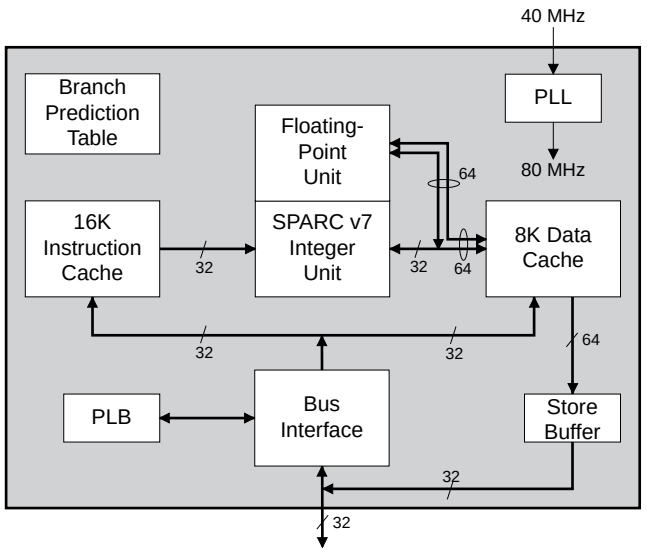

Figure 2. The SPARC Power (SPX2) microprocessor contains a SPARC CPU, FPU, and 24K of cache on a single chip.

Operation	SPARC2 (W8701) at 40 MHz		SPARC Power at 80 MHz	
Integer Load Single	2 cycles	50 ns	12 ns	1 cycle
Integer Load Double	3 cycles	75 ns	25 ns	2 cycles
Integer Store Single	3 cycles	75 ns	25 ns	2 cycles
Integer Store Double	4 cycles	100 ns	37 ns	3 cycles
FP Load Single	2 cycles	50 ns	12 ns	1 cycle
FP Load Double	3 cycles	75 ns	12 ns	1 cycle
FP Store Single	3 cycles	75 ns	25 ns	2 cycles
FP Store Double	4 cycles	100 ns	25 ns	2 cycles
FP ALU Single	4 cycles	100 ns	50 ns	4 cycles
FP ALU Double	4 cycles	100 ns	50 ns	4 cycles
FP Multiply Single	4 cycles	100 ns	50 ns	4 cycles
FP Multiply Double	6 cycles	150 ns	75 ns	6 cycles
FP Divide Single	22 cycles	550 ns	425 ns	34 cycles
FP Divide Double	36 cycles	900 ns	775 ns	62 cycles
FP Sq. Root Single	33 cycles	825 ns	425 ns	34 cycles
FP Sq. Root Double	62 cycles	1550 ns	775 ns	62 cycles

Table 1. SPARC Power cycle counts for many operations have been reduced from the 8701, resulting in greater than 2× improvement in these functions when measured in nanoseconds.

Doubled Speed With Same IC Process

Figure 2 shows Weitek's solution to these problems. Maintaining compatibility with the 8701, the SPX2 contains a complete SPARC (version 7) integer unit and floating-point unit (FPU). Since both the 8701 and SPX2 use the same 0.8-micron, three-layer-metal CMOS process, these units had to be completely redesigned to reach the 80-MHz clock speed.

Impressively, the SPX2 does not resort to superpipelining or half-micron CMOS to reach this speed. In fact, Weitek's chip has the fastest clock rate of any available SPARC processor and, except for HP's PA7100, is the fastest non-superpipelined CPU of any variety. The SPX2 uses a classic five-stage pipeline—instruction-cache access, decode, execute, data-cache access, and writeback—to keep the design simple. Weitek also eschewed the complexities of a superscalar design, choosing instead a single-issue model.

Since the team could not count on any increase in circuit speed, the pipeline had to be carefully designed to meet the 80-MHz goal. A critical path in the 8701 occurs when an instruction modifying the condition codes is immediately followed by a conditional branch. In this case, the register read and ALU operation take place during the decode stage; during the execute stage, the condition codes are evaluated and the selected instruction address must be driven off the chip to the external cache with enough margin to meet the address set-up time.

In the SPX2, the lengthy register read and ALU operations are split into separate stages by moving the ALU into the execute stage. Even with the on-chip cache, however, there is too little time to evaluate the condition codes and select the correct instruction address by the end of the execute stage.

The SPX2 uses branch prediction to overcome this problem. It employs a 256-entry branch history table, accessed during the decode stage, to predict the outcome of each conditional branch using the four-state Lee and Smith algorithm *(see **070402.PDF**)* also used by Pentium. If the branch is correctly predicted, there is adequate time to fetch the next instruction with no branch penalty. On an incorrect prediction, there is a single-cycle penalty (after the delay-slot instruction).

The 8701 uses a four-stage pipe with the writeback stage immediately following the execute stage. This creates a critical path in resolving all possible exceptions by the end of the execute stage. The SPX2 adds the "data-cache access" stage to ease these restrictions. A separate 32-bit adder for load-address generation simplifies the control logic for loads.

Floating-point math has been Weitek's specialty, and the SPX2's floating-point unit doubles the performance of the 8701 by maintaining the same cycle counts for basic FP operations, as shown in Table 1. The 8701 also requires a "dead" cycle between FP operations that is eliminated in the SPX2, further increasing performance. Like the 8701, the new FPU contains two independent units, one for multiplication and one for all other operations; the two units can execute in parallel, and integer instructions can also execute while the FPU is in use.

On-Chip Caches Speed Loads and Stores

The memory architecture has been significantly improved over the 8701, which uses a single 32-bit path to its external cache. The SPX2 implements separate 32-bit buses for data and instructions, eliminating the pipeline stall caused by loads and stores in the 8701. The on-chip data cache can provide 64 bits to the FPU in a single

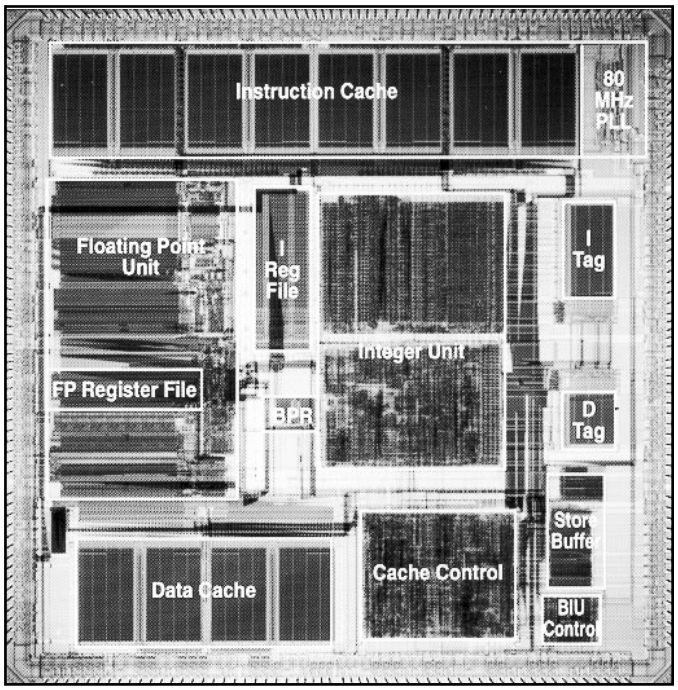

Figure 3. Die photo of the SPARC Power chip, which uses 1.8 million transistors on a 16.1 mm × 16.2 mm (634 × 638 mils) die.

cycle. These changes increase the peak cache bandwidth by a factor of six and greatly reduce the time needed to perform loads and stores, as shown in Table 1.

Unlike most modern microprocessors, the SPX2 must work with an off-chip MMU (memory-management unit). Thus, the on-chip caches are indexed and tagged with virtual addresses, as is the external cache. This allows address translation to occur in parallel with the cache access. To reduce complexity, the on-chip caches are direct-mapped and use the same 32-byte line size as the external cache.

To maintain compatibility with existing software, the on-chip caches must be totally transparent. The external cache is kept up-to-date via a write-through protocol. All flush operations intended for the external cache are shadowed in the on-chip caches. Thus, the CPU does not have to respond to any cache snoops.

On a read miss in the primary cache, the data is requested from the secondary cache. If this request also misses, the external cache controller processes the miss and updates the external cache. The SPX2 snoops the data as it comes across the processor bus and updates its internal caches. The cache controller fetches the "critical" word first, and the CPU begins execution as soon as that word is received. The SPX2 also implements "streaming," where a stream of consecutive words (typically instructions) can be taken directly from the processor bus on subsequent cycles, so long as there are no pipeline stalls.

On a write miss, the selected cache line must be invalidated, even if it has a different virtual tag, because it could be mapped to the same physical address. Each line in the on-chip data cache is divided into two 16-byte sub-blocks. Each sub-block has its own valid bit so that a write miss invalidates only half of a primary cache line, increasing the cache hit ratio by allowing continued access to the other half. An invalidated sub-block can later be filled from the external cache without updating the entire line.

Since the write-through data cache generates a considerable amount of write traffic, a four-entry store buffer is used to allow the CPU to proceed after posting a store. Each entry can contain either a 32-bit (single word) or 64-bit (double word) value. To avoid stalling writes that the MMU indicates are non-cacheable, the store buffer only posts writes that have hit in the external cache. The SPX2 includes a protection look-aside buffer (PLB), which keeps track of the four cacheable pages that were most recently written, to reduce accesses to the external MMU. The PLB is a subset of the TLB entries in that MMU.

Price and Availability

The SPARC Power upgrade kit is currently in beta test, with production shipments available by September. The cost of the upgrade kit is $1500 for single-unit orders. For more information, contact Weitek at 1060 East Arques Avenue, Sunnyvale, CA 94086; 408/738-8400. For Faxback information, call 800/827-8708.

Big Chip in a Small Package

Figure 3 shows a die photo of the SPX2. The processor uses 1.8 million transistors, about six times as many as the 8701. Most (1.3 million transistors) of the increase comes from the 24K of on-chip cache. The integer unit has increased in size from about 50,000 to 150,000 transistors due to the branch history table, load-address adder, and more complex control logic. The floating-point units in both chips use about 150,000 transistors each. The remaining transistors on the SPX2 are used for cache control and the bus interface.

The die is quite large at 261 mm^2, bigger than even SuperSPARC+ (256 mm^2). The SPX2 uses a less expensive manufacturing process than the TI chip's 0.7-micron BiCMOS process, so Weitek will have somewhat lower manufacturing costs. The SPX2 is packaged in a 207-pin PGA similar to the 8701.

At a clock rate of 80 MHz, the SPX2 has a maximum power dissipation of 6 watts. This is a bit greater than the 4.5 watts of the 8701 but is low compared to Pentium, Alpha, and other contemporary microprocessors. The SPX2 comes with a small (0.2-inch tall) heat sink to ensure adequate cooling in an SS2 or IPX system.

More Than 40% Performance Gain

Table 2 compares the performance of identical SPEC binaries run on SS2, upgraded SS2, and low-end SS10 systems. For comparison, a high-end SS10 running optimized code is included. This data shows that the SPX2 provides a 51% improvement in SPECint92 and a 42% improvement on SPECfp92 over an 8701. Weitek claims that most applications will see a similar 40%–50% increase, with a few CPU-intensive programs in the 60%–90% range. Although the core CPU provides a 2× speedup, the small on-chip caches and small (by today's standards) 64K secondary cache cannot fully compensate for the doubling (in CPU cycles) of the main memory latency.

Weitek's figures for the SS2 are slightly lower than Sun's official SPEC ratings, which are 21.8 on integer and 22.8 on floating-point. Weitek was unable to duplicate Sun's figures due to compiler differences, but both sets of figures are fairly close. The SS10/30 measurements are about 20%–30% lower than Sun's figures for that system (45 SPECint92 and 54 SPECfp92); Weitek claims that this difference represents the performance penalty for running unoptimized code on a SuperSPARC system, which would be the case for users who have bought a new system without upgrading their software. In this scenario, the SPX2 is actually close to the performance of a cacheless, 36-MHz SS10 system.

The SPX2 is blown away, however, by a full-speed SuperSPARC+ running recompiled code. The TI chip uses superior IC technology to incorporate a superscalar integer unit as well as an extra 12K of on-chip cache. Another difference, however, is that the SPX2 must function within the antiquated SS2 system design, with its external MMU and only 64K of external cache; the SS10, by comparison, uses 1M of cache to achieve its superior performance. Because of the design constraints on the upgrade chip, it is not really fair to compare its performance to a processor designed from scratch.

System	SPARC Station 2	SPARC Station 2	SPARC Station 10/30	SPARC Station 10/51
Processor	SPARC Power	SPARC2 (W8701)	TI Super	TI Super+
Clock Rate	80 MHz	40 MHz	36 MHz	50 MHz
Cache (on/off-chip)	24K/64K	none/64K	36K/none	36K/1M
espresso	35.5	18.5	32.7	58.7
li	32.6	23.2	38.5	63.3
equtott	34.8	21.8	37.4	108.9
compress	22.6	17.3	26.2	38.2
sc	42.9	28.5	50.4	93.0
gcc	28.3	21.0	36.3	53.4
SPECint92	32.2	21.4	36.2	65.2
spice	22.8	17.5	27.8	56.2
doduc	24.9	17.9	38.7	81.2
mdljdp2	45.5	27.3	52.2	86.4
wave5	22.7	15.6	24.7	57.9
tomcatv	22.2	17.3	36.2	77.3
ora	60.9	44.1	104.8	160.3
alvinn	47.4	34.3	68.7	175.6
ear	49.9	26.9	52.4	95.0
mdljsp2	28.6	15.8	26.3	42.2
swm256	18.8	13.8	27.5	43.7
su2cor	30.2	27.3	31.0	113.3
hydro2d	29.2	19.9	40.7	79.1
nasa7	26.4	22.5	30.9	93.5
fpppp	34.1	21.7	43.0	101.7
SPECfp92	31.1	21.8	39.6	83.0

Table 2. Running identical binaries, SPARC Power is significantly faster than SS2 and comes close to a low-end SS10, but falls behind the high-end SuperSPARC+. (Source: Weitek)

Niche Marketing

The SPARC Power upgrade processor is designed for a single, short-term market niche. Although this niche is small by PC standards, the potential opportunity to upgrade a quarter of a million systems would be attractive to most RISC vendors. Weitek, in particular, seized this opportunity as what is probably its last chance to leverage its investment in the SPARC microprocessor business. TI and Fujitsu have already been selected to build the follow-ons to SuperSPARC+ and microSPARC, and Fujitsu has the inside track for Sun's next-generation high-end chip, UltraSPARC.

The SPX2 is a good solution for its target market. By cleverly arranging the pipeline design and taking advantage of on-chip primary caches, Weitek was able to double the internal clock speed of the chip without moving to a more expensive IC technology. Even with its large die size, the design should provide outstanding gross margins at a $1500 unit price. Although this seems expensive for a single chip, it is a cost-effective alternative to a $16,000 workstation and associated software.

An important goal of the program was to minimize design costs. Given the uncertain sales situation, Weitek chose to limit performance to the minimum needed to offer a significant upgrade. Enhancements such as superscalar dispatch or set-associative caches could have provided a greater boost in performance but would have significantly added to both the development time and the design costs.

Intel's experience with its OverDrive program is that, so far, few users have chosen to upgrade. While it is too early to count out OverDrive, Weitek has a few advantages over Intel. Sun's customer base tends to be more technically savvy and willing to consider a processor upgrade. More importantly, there is pent-up dissatisfaction with Solaris 2.0 that could be vented by sticking with a processor compatible with SunOS. Many users, however, will choose to upgrade to a new SPARCstation 10 for its higher performance and broader range of processor upgrades. Weitek hopes that 20% or even 40% of SS2 and IPX users choose its upgrade chip, yielding more than enough revenue to cover the limited costs of the SPARC Power program. ♦

